Logical Data Independence in the 21st Century – Co-Existing Schema Version Support with InVerDa

Kai Herrmann
Technische Universität
Dresden, Germany
kai.herrmann@tu-dresden.de

Hannes Voigt
Technische Universität
Dresden, Germany
hannes.voigt@tu-dresden.de

Andreas Behrend
Universität Bonn
Germany
behrend@iai.uni-bonn.de

Jonas Rausch
Technische Universität
Dresden, Germany
jonas.rausch@tu-dresden.de

Wolfgang Lehner
Technische Universität
Dresden, Germany
wolfgang.lehner@tu-dresden.de

arXiv:1608.05564v1 [cs.DB] 19 Aug 2016

ABSTRACT

We present InVerDa, a tool for end-to-end support of co-existing schema versions within one database. While it is state of the art to run multiple versions of a continuously developed application concurrently, the same is hard for databases. In order to keep multiple co-existing schema versions alive—that all access the same data set—developers usually employ handwritten delta code (e.g. views and triggers in SQL). This delta code is hard to write and hard to maintain: if a database administrator decides to adapt the physical table schema, all handwritten delta code needs to be adapted as well, which is expensive and error-prone in practice. With InVerDa, developers use a simple bidirectional database evolution language in the first place that carries enough information to generate all the delta code automatically. Without additional effort, new schema versions become immediately accessible and data changes in any version are visible in all schema versions at the same time. We formally validate the correctness of this propagation. InVerDa also allows for easily changing the physical table designs without affecting the availability of co-existing schema versions. This greatly increases robustness (264 times less lines of code) and allows for significant performance optimization.

1. INTRODUCTION

Database management systems (DBMSs) lack proper support for co-existing schema versions within the same database. With today's realities in information system development and deployment—namely agile development methods, code refactoring, short release cycles, stepwise deployment, varying update adoption time, legacy system support, etc.—such support becomes increasingly desirable. Development tools such as GIT, SVN, and Maven make it feasible to maintain multiple versions of the same application and deploy and run several of these versions concurrently. The same is hard for a database, though. Databases live in complex, deeply inter-connected worlds [3]. In enterprise information systems databases feed hundreds of subsystems across the company domain, connecting decades-old legacy systems with brand new web front ends or innovative analytics pipelines. All of these subsystems are typically run by different stakeholders, with different development cycle, and different upgrade constraints. Typically, adopting changes to a database schema—if even possible on the client-side—will stretch out over time. Hence, the database schema versions these subsystems expect have to be kept alive on the database side so that the database continues to serve them. Databases with co-existing schema versions, i.e. databases that are exposed through multiple schemas, are a reality in large information systems developers and DBAs have to cope with.

Unfortunately, current database management systems (DBMSs) lack proper support for co-existing schema versions. Current DBMSs essentially force developers to migrate a database completely in one haul to a new schema version. Keeping older schema versions alive to continue serving all database clients independently from their adoption time is costly. Before and after a one-haul migration it requires manually written and maintained delta code either in the database or in the application. Alternatively, a database can be kept in multiple replicated instances each in a different schema version. Although this makes reading each schema version very simple, various ELT jobs (a.k.a. delta code) have to be developed and maintained for propagation of updates to prevent inconsistencies. Either way co-existing schema versions require considerable development resources. In fact, data migration accounted on average for 31 % of IT project budgets in 2011 according to [10] and co-existing schema versions play a considerable role in that. In short, handling co-existing schema versions is very costly and error-prone and forces developers into less agility, longer release cycles, riskier big-bang migration, etc.

We envision a DBMS-integrated co-existing schema version support allowing a single database to have multiple co-existing schema versions that all can be accessed at the same time. Developers can evolve an existing database schema with simple but powerful schema modification operations (SMOs) to add a new schema version to the database. New schema versions become immediately available. Applications can read and write data through any schema version concurrently; writes in one version are reflected in all other versions. Old schema versions can be easily dropped without losing any data that is visible in other schema versions. Furthermore, the database administrator (DBA) can easily configure in which schema version the data is primarily materialized. If the workload mix changes, the DBA can change the materialization without affecting the availability of any schema version and without a developer being involved. DBMS-integrated co-existing schema version support completely eliminates the need to write a single line of delta code.

Valuable tools, such as Liquibase[1], Rails Migrations[2], and DBmaestro Teamwork[3], help to manage schema versions outside the DBMS and generate SQL scripts for migrating to a new schema version. They mitigate the data migration cost, but their focus is on schema evolution. Support for co-existing schemas is very limited. The inspiring research work PRIMA [11] considers co-existence of schema versions a bit but only in the sense of delayed propagation of data from an old to a newer version. PRIMA offers forward propagation of updates and backward propagation of reads, but not vice versa. To our best knowledge, there is no full-fledged support for co-existing schema versions in DBMSs yet.

In this paper, we present INVERDA (Integrated Versioning of Database schemas), a concept and an implementation to extend DBMSs with the envisioned co-existing schema versions. For the developer, INVERDA provides the Bidirectional Database Evolution Language BIDEL. BIDEL has an evolution command and SMOs. With them the developer can easily evolve a database schema by adding and deleting columns, partitioning, and merging tables, etc. In that respect, BIDEL is of equal power and intuitiveness as established database evolution languages (DELs) [4, 8]. BIDEL's unique feature, however, is the bidirectionality of its SMOs, the crucial ingredient to full-fledged support for co-existing schema versions. Bidirectional SMOs facilitate the propagation of reads and writes in any direction between schema versions. More precisely, upon an evolution command, INVERDA generates all necessary delta code in the database, i.e. the complete access propagation logic between all schema versions is implemented with the click of a button. The delta code generation is very fast and allows INVERDA to make new schema version immediately available upon submission of an evolution command. INVERDA greatly simplifies evolution tasks, since the developer only has to write a BIDEL script and no delta code at all.

For the DBA INVERDA provides a very compact and declarative database migration command. It allows the DBA to easily change the physical data representation into a different schema version with just a single line of code. INVERDA automatically migrates all affected data and regenerates all delta code—no interaction with the developer required. This opens opportunities for physical optimization if the workload focus shifts between schema versions because, e.g., the majority of database client applications has been upgraded to a new version. Without INVERDA such an optimization requires to rewrite all affected delta code manually. The main contributions of INVERDA are:

System integration To the best of our knowledge, we are the first to present end-to-end support for co-existing schema versions. Providing developers and DBAs with an intuitive interface to write and execute evolution/migration scripts, INVERDA automatically generates all required delta code to keep all schema versions alive. We show how to integrate the INVERDA concept into an existing DBMSs by reusing existing database functionality as much as possible. (Section 3)

Bidirectional database evolution language BIDEL is a DEL with bidirectional SMOs for translating data accesses between schema versions. We introduce the BIDEL SMOs and formally validate the correct and information preserving bidirectionality of BIDEL's SMOs. In the evaluation, we show that BIDEL requires significantly (up to 264 x) less code than evolutions and migrations manually written in SQL. (Sections 4 and 5)

Co-existing schema versions INVERDA generates delta code from BIDEL-specified schema evolutions to allow reads and writes on all co-existing schema versions, each providing an individual view on the same shared dataset. Each schema version is a full-fledged database schema, which transparently hides the versioning from the user. We describe the delta code generation procedure. In the evaluation, we show that delta code generation is really fast (< 1 s) and the generated delta code shows query performance comparable to hand-written delta code. (Section 6)

Decoupled materialization INVERDA makes manual migrations obsolete by providing simple declarative migration commands. Such a single-line migration command triggers the physical data movement as well as the adaptation of all involved delta code. It allows the DBA to optimize the physical table schema of the database independently from the schema evolution. We introduce the database migration command and describe the migration procedure. In the evaluation, we show that the physical optimizations INVERDA allows with its migration command easily compensate for the overhead it implies. (Section 7)

The paper focuses on the principles and concepts INVERDA is built on. In this respect, INVERDA is not linked to any specific DBMS but rather serves as a blue print to implement support for co-existing schema versions in any relational DBMS. We present a prototypical integration of INVERDA which realizes the delta code as views and triggers in a common DBMS. Note, the INVERDA concept is not limited to generating views and triggers. Alternatively, INVERDA could generate delta code e.g. within the application or as optimized execution plans within the database. The implementation presented in this work is merely one possibility which allows us to focus on INVERDA's conceptual strength: generating delta code for co-existing schema versions from a bidirectional database evolution language.

[1] http://www.liquibase.org/
[2] http://guides.rubyonrails.org/migrations.html
[3] http://www.dbmaestro.com/

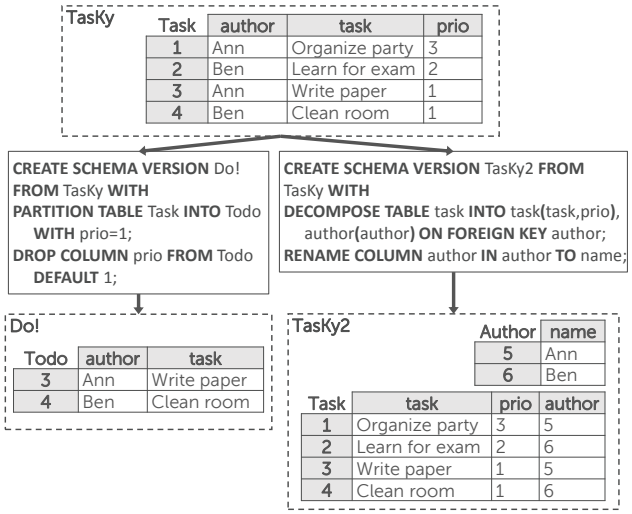
TasKy
Task author task prio
1 Ann Organize party 3
2 Ben Learn for exam 2
3 Ann Write paper 1
4 Ben Clean room 1
CREATE SCHEMA VERSION Do! FROM TasKy WITH
PARTITION TABLE Task INTO Todo WITH prio=1;
DROP COLUMN prio FROM Todo DEFAULT 1;
CREATE SCHEMA VERSION TasKy2 FROM TasKy WITH
DECOMPOSE TABLE task INTO task(task,prio), author(author) ON FOREIGN KEY author;
RENAME COLUMN author IN author TO name;
Do!
Todo author task
3 Ann Write paper
4 Ben Clean room
TasKy2
Author name
5 Ann
6 Ben
Task task prio author
1 Organize party 3 5
2 Learn for exam 2 6
3 Write paper 1 5
4 Clean room 1 6

Figure 1: Example.

After introducing InVerDa from a user perspective in Section 2, we present the basic architecture of InVerDa in Section 3. In Section 4, we present BiDEL in detail and formally validate its bidirectionality in Section 5. In Sections 6 and 7 we sketch how to generate delta code and how to change the materialization schema. Finally, we evaluate InVerDa in Section 8, discuss related work in Section 9, and conclude the paper in Section 10.

2. USER PERSPECTIVE ON INVERDA

In the following, we introduce the co-existing schema version support by InVerDa using the example of a simple task management system called TasKy (cf. Figure 1). TasKy is a desktop application which is backed by a central database. It allows users for creating new tasks as well as listing, updating, and deleting them. Each task has an author and a priority. Tasks with priority 1 are most urgent ones. The first release of TasKy stores all its data in a single table `Task(author,task,prio)`. TasKy has productive go-live and users begin to feed the database with their tasks.

Developer. TasKy gets widely accepted and after some weeks it is extended by a third party phone app called Do! to list most urgent tasks. However Do! expects a different database schema than TasKy is using. The Do! schema consists of a table `Todo(author,task)` containing only tasks of priority 1. Obviously, the initial schema version needs to stay alive for TasKy, which is broadly installed. InVerDa greatly simplifies the job as it handles all the necessary delta code for the developer. The developer merely executes the BiDEL script for Do! as shown in Figure 1, which instructs InVerDa to derive schema Do! from schema TasKy by creating a horizontal partition of `Task` with `prio=1` and dropping the priority column. Executing the script creates a new schema including the view `Todo` as well as delta code for propagating data changes. When a user inserts a new todo into `Todo`, this will automatically insert a corresponding task with priority 1 to `Task` in TasKy. Equally, updates and deletions are propagated back to the TasKy schema. Hence, the TasKy data is immediately available to be read and written through the newly incorporated Do! app by simply executing the above BiDEL evolution script.

Assume the TasKy application is further refined to make it more useful for its users. For the next release TasKy2, it is decided to normalize the table `Task` into `Task` and `Author`. For a stepwise roll-out of TasKy2, the old schema of TasKy has to remain alive until all clients have been updated. Again, InVerDa does the job. When executing the BiDEL script as shown in Figure 1, InVerDa creates the schema version TasKy2 and decomposes the ConTab `Task` to separate the tasks from their authors while creating a foreign key to maintain the dependency. Additionally, the column `author` is renamed to `name`. InVerDa generates delta code to make the TasKy2 schema immediately available. Write operations to any of the three schema versions are propagated to all other schema versions.

DBA. The physical tables initially used for storing the data are determined by create table statements. All other schema versions are implemented with the help of delta code. The delta code introduces an overhead on read and write accesses to new schema versions. The more SMOs are between schema versions, the more delta code is involved and the higher is the overhead. In the case of the task management system, the schema versions TasKy2 and Do! have delta code towards the physical table `Task`. Assume, some weeks after releasing TasKy2 the majority of the users has upgraded to the new version. TasKy2 comes with its own phone app, so that the schemas TasKy and Do! are still accessed but merely by a minority of users. Hence, it seems appropriate to migrate data physically to the table version of the TasKy2 schema, now. Traditionally, developers would have to write a migration script, which moves the data and implements new delta code. All that can accumulate to some hundred lines of code, which need to be tested intensively in order to prevent it from messing up our data. With InVerDa, this task can be solved with a single line:

```
MATERIALIZE 'TasKy2';
```

Upon this statement, InVerDa transparently runs the physical data migration to schema TasKy2, maintaining transaction guarantees, and updates the involved delta code of all schema versions. No developers need to be involved. All schema versions stay available; read and write operations are merely propagated to a different set of physical tables, now. In sum, InVerDa allows the user to continuously use all schema versions, the developer to continuously developing the applications, and the DBA to independently adapting the physical table schema to the workload.

3. INVERDA ARCHITECTURE

InVerDa can be simply integrated into an existing architectures of a relational DBMS. It adds schema evolution functionality and support for co-existing schema versions. InVerDa functionality is exposed to users via two interfaces: (1) BiDEL (bidirectional database evolution language) and (2) migration commands.

BiDEL provides a comprehensive set of bidirectional SMOs in order to create a new schema version either starting from an empty schema or as an evolution from a given schema version. SMOs allow to create or drop or rename tables and columns, (vertically) decompose or join tables, and (horizontally) partition or merge tables (cf. Figure 2). This established set of SMOs has already proved to be practically and theoretically comprehensive [4, 8]. The unique feature of BiDEL SMOs is bidirectionality. Essentially, the

```
CREATE SCHEMA VERSION name_new [FROM name_old]
       WITH SMO_1; ... SMO_n;
DROP SCHEMA VERSION version_n;
- - - - - - - - - - - - - - - - - - - - - - - - - - - - -
CREATE TABLE R(c_1,...,c_n)
DROP TABLE R
RENAME TABLE R INTO R'
RENAME COLUMN r IN R_i TO r'
ADD COLUMN a AS f(r_1,...,r_n) INTO R_i
DROP COLUMN r FROM R_i DEFAULT f(r_1,...,r_n)
DECOMPOSE TABLE R INTO S(s_1,...,s_n)
          [, T(t_1,...,t_m) ON (PK|FK fk|cond)]
[OUTER] JOIN TABLE R, S INTO T ON (PK|FK fk|cond)
PARTITION TABLE T INTO R WITH c_R [, S WITH c_S ]
MERGE TABLE R (c_R), S (c_S) INTO T
```

Figure 2: BiDEL SMOs.

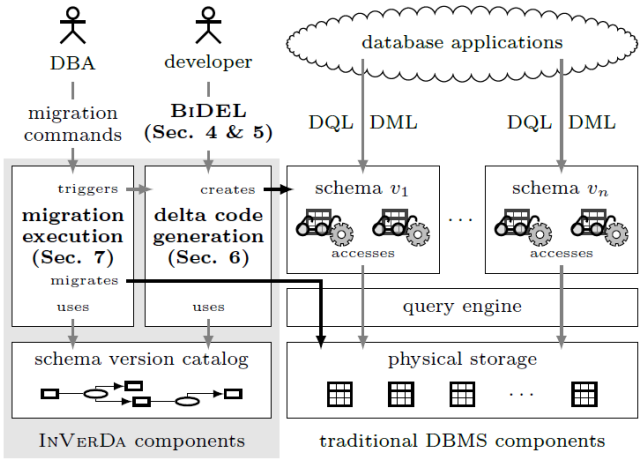

Figure 3: InVerDa integration into DBMS.

arguments of each BiDEL SMO gather enough information to facilitate full propagation of reads and writes between schema versions in both directions, forward propagation from the old to the new version as well as backward propagation from the new to the old version. For instance, `DROP COLUMN` requires a function $f(r_1, \ldots, r_n)$ that computes the value for the dropped column if a tuple, inserted in the new schema version, is propagated back to an old schema version. Finally, BiDEL allows dropping old schema versions that are not used any longer, which merely drops the schema version itself but maintains the data if still is still needed in other versions.

InVerDa's migration commands allow for changing the physical data representation. By default data is materialized in the source schema version. Assuming a table is partitioned in a schema evolution step, the data remains physically unpartitioned. With a migration command the physical data representation of this SMO instance can be changed so that the data is also physically partitioned—the SMO instance is called materialized then. Within this work, we focus on non-redundant materialization which means the data is stored either on the source or the target side of the SMO but not on both. Migration commands are very simple. They either materialize a set of table versions or a complete schema version. The latter is merely a convenience command allowing to materialize multiple table versions in one step.

InVerDa utilizes existing DBMS components such as physical data storage, indexing, transaction handling, query processing, etc. In our prototypical implementation, the delta code itself consists of standard views and triggers. Data is stored in regular tables. Hence, all processing of DQL/DML commands requested by database applications happens in the standard query engine of the DBMS.

So, InVerDa's components are only active during database development and migrations. In order to process data accesses of database applications at runtime, merely the generated views and triggers are used and no InVerDa components are involved. Note that the employed triggers can lead to a cascaded execution but in a controlled manner as there are no cycles in the version genealogy.

Figure 3 outlines the principle components of an InVerDa-equipped DBMS. As can be seen, InVerDa adds three components to the DBMS: (1) the *delta code generation* creates views and triggers to expose schema versions based on the current physical data representation. Delta code generation is either triggered by a developer issuing BiDEL commands to create a new schema version or by the DBA issuing migration commands to change the physical data representation. The delta code itself are standard commands of the DBMS's query engine. (2) the *migration execution* orchestrates the actual migration of data from one physical representation to another and the adaptation of the delta code. The data migration is done with the help of query engine capabilities. (3) the *schema version catalog* maintains the genealogy of schema versions. It represents with which SMOs tables of a schema version have been evolved to tables of another schema version as well as in which schema version the data is actually stored physically. Figure 4 shows the schema version catalog for the example from Figure 1 with the initial materialization (all SMOs virtualized).

When developers execute BiDEL scripts, the respective SMO instances and table versions are registered in the schema version catalog. The schema version catalog maintains references to tables in the physical storage that hold the payload data and to auxiliary tables that hold otherwise lost information of the not necessarily information preserving SMOs. The materialization state of the SMOs, which can be changed by the DBA through migration commands, determine which data tables and auxiliary tables are physically present or not. InVerDa uses the schema version catalog to generate delta code for new schema versions or for changed physical table schemas. Data accesses of database applications are processed by the generated delta code within the DBMS's query engine. When a developer drops an old schema version that is not used any longer, the schema version is removed from the catalog. However, the respective SMOs are only removed from the catalog as well if they are no longer part of an evolution that connects two remaining schema versions.

The schema version catalog is the central knowledge base and documentation for all created schema versions and the evolution between them. To this end, the catalog stores the genealogy of schema versions by means of a directed acyclic hypergraph (T, E). Each vertex $t \in T$ of that graph represents a table version. Each hyperedge $e \in E$ represents one SMO instance, i.e., one table evolution step. An SMO instance $e = (S, T)$ evolves a set of source table versions S into a set of target table versions T. Additionally, the schema version catalog stores for every SMO instance the

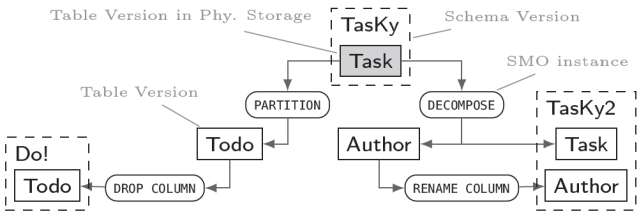

Figure 4: Example of schema version catalog.

SMO type (partition, merge, etc.), the parameter set, and its state of materialization. Each schema version is basically a subset of all table versions in the system. Schema versions share a table version if that table has not been evolved from one schema version to another. At evolution time, InVerDa uses this information to generate delta code that makes all schema versions accessible. At query time, the generated delta code itself is executed by the existing DBMS's query engine—outside InVerDa components.

4. BIDEL - BIDIRECTIONAL SMOS

BiDEL's unique feature is the bidirectional semantics of its SMOs. This property is the key feature which forms the basis for InVerDa's new functionalities. We highlight the design principles behind BiDEL SMOs and formally validate their bidirectionality. All BiDEL SMOs follow the same design principles. Without loss of generality, the `PARTITION` SMO is used as a representative example in this section.

Figure 5 illustrates the principle setup of a single SMO instance resulting from the sample statement

`PARTITION TABLE` T **`INTO`** R **`WITH`** c_R**`,`** S **`WITH`** c_S

which horizontally partitions a table T from the source schema into two tables R and S in the target schema based on conditions c_R and c_S. Assuming both schemas are materialized, reads and writes on both schema version can simply be delegated to the corresponding data tables T_D, R_D, and S_D, respectively. However, InVerDa materializes data non-redundantly on one side of the SMO instance, only. If the data is physically stored on the source side of an SMO instance, the SMO instance is said to be *virtualized*; with data stored on the target side it is said to be *materialized*. In any case, reads and writes on the unmaterialized side have to be mapped to the materialized side.

The semantics of each SMO is defined by two functions γ_{trg} and γ_{src} which describe precisely the mapping from the source side to the target side and vice versa, respectively. Assuming the target side of `PARTITION` is materialized, all reads on T are mapped by γ_{src} to reads on R_D and S_D; and writes on T are mapped by γ_{trg} to writes on R_D and S_D. While the payload data of R, S, and T is stored in the physical tables R_D, S_D, and T_D, the tables R^-, S^+, S^-, R^*, S^*, and T' are auxiliary tables that `PARTITION` needs for the bidirectional propagation. Note that the semantics of SMOs is complete if reads and writes on both, source and target schema work correctly without any loss or gain of information regardless on which of both sides the data is physically stored.

For the definition of γ_{trg} and γ_{src}, we use Datalog rules—more precisely Datalog rules templates that get instantiated with the parameters of an SMO instance. For brevity of presentation we use some extensions to the standard Datalog

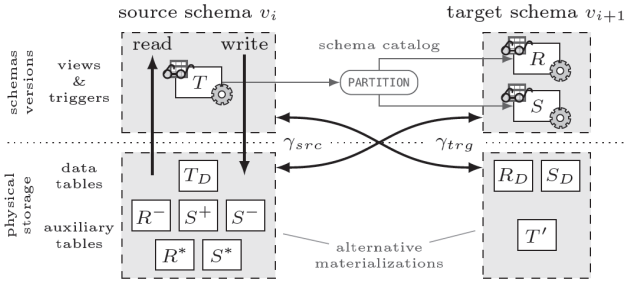

Figure 5: Mapping functions of single `PARTITION` SMO.

syntax: For variables, small letters denote single attributes and capital letters lists of attributes. For equality predicates on attribute lists, both lists need to have the same length and same content, i.e. for $A = (a_1, \ldots, a_n)$ and $B = (b_1, \ldots, b_m)$, $A = B$ holds if $n = m \wedge a_1 = b_1 \wedge \ldots \wedge a_n = b_n$. All tables have an attribute p, which is an InVerDa-managed identifier for each tuple. The key p is necessary to uniquely identify any tuple across versions. Additionally, p ensures that the multiset semantics of a relational database fits with the set semantics of Datalog, as the unique key p prevents equal tuples in one relation. For a table T we assume $T(p, \_)$ and $\neg T(p, \_)$ to be save predicates since any table can be projected to its key.

For the exemplary `PARTITION` instance, let's assume this instance is materialized, i.e. data is stored on the target side, and let's consider the γ_{trg} mapping function first. `PARTITION` horizontally splits a table T from the source schema into two tables R and S in the target schema based on conditions c_R and c_S:

$$R(p, A) \leftarrow T(p, A), c_R(A) \quad (1)$$
$$S(p, A) \leftarrow T(p, A), c_S(A) \quad (2)$$

The conditions c_R and c_S can be arbitrarily set by the user so that Rule 1 and Rule 2 are insufficient for achieving the desired bidirectional semantics. This is caused by the fact that the source table T may contain tuples neither captured by c_R nor by c_S. Such tuples are stored on the target side in the auxiliary table T':

$$T'(p, A) \leftarrow T(p, A), \neg c_R(A), \neg c_S(A) \quad (3)$$

Let's now consider the γ_{src} mapping function for reconstructing T while the target side is still considered to be materialized. Reconstructing T from the target side is essentially a union of R, S, and T'. Nevertheless, c_R and c_S are not necessarily disjoint. One source tuple may occur as two equal but independent instances in R and S. We call such two instances *twins*. Twins can be updated independently resulting in *separated twins*, i.e. two tuples—one in R and one in S—with equal key p but different in the other attribute values. In order to resolve this ambiguity, we consider the first twin in R to be the primus inter pares and hence define γ_{src} of `PARTITION` to propagate back all tuples in R as well as those tuples in S not contained in R:

$$T(p, A) \leftarrow R(p, A) \quad (4)$$
$$T(p, A) \leftarrow S(p, A), \neg R(p, \_) \quad (5)$$
$$T(p, A) \leftarrow T'(p, A) \quad (6)$$

The Rules 1–6 define sufficient semantics for `PARTITION` as long as the target side is materialized.

Let's now assume the SMO instance is virtualized, i.e. data is stored on the source side, and let's keep considering the γ_{src} mapping function. Again, R and S can contain separated twins—unequal tuples with equal key p. According to Rule 5, T stores only the separated twin from R. To avoid losing the other twin in S, it is stored in the auxiliary table S^+:

$$S^+(p, A) \leftarrow S(p, A), R(p, A'), A \neq A' \tag{7}$$

Accordingly, γ_{trg} has to reconstruct the separated twin in S from S^+ instead of T (concerns Rule 2). Twins can also be deleted independently resulting in a *lost twin*. Given the data is materialized on the source side, a lost twin would be directly recreated from its other twin via T. To prevent this information gain and keep lost twins lost, γ_{src} keeps the keys of lost twins from R and S in auxiliary tables R^- and S^-:

$$R^-(p) \leftarrow S(p, A), \neg R(p, \_), c_R(A) \tag{8}$$

$$S^-(p) \leftarrow R(p, A), \neg S(p, \_), c_S(A) \tag{9}$$

Accordingly, γ_{trg} has to exclude lost twins stored in R^- from R (concerns Rule 1) and those in S^- from S (concerns Rule 2). Twins result from data changes issued to the target schema containing R and S which can also lead to tuples that do not meet the conditions c_R resp. c_S. In order to ensure that the reconstruction of such tuples is possible from a materialized table T, auxiliary tables R^* and S^* are employed for identifying those tuples using their identifiers. Note that this does not apply for tuples which have been already evacuated to S^+. According to Rule 7, this means a corresponding identifier is either not in R or the data A in R is equal to the data in S.

$$S^*(p) \leftarrow S(p, A), \neg c_S(A), \neg R(p, \_) \tag{10}$$

$$S^*(p) \leftarrow S(p, A), \neg c_S(A), R(p, A) \tag{11}$$

$$R^*(p) \leftarrow R(p, A), \neg c_R(A) \tag{12}$$

The complete sets of rules for defining γ_{trg} respectively γ_{src} are defined as follows:

$\gamma_{\mathbf{trg}}$:

$$R(p, A) \leftarrow T(p, A), c_R(A), \neg R^-(p) \tag{13}$$

$$R(p, A) \leftarrow T(p, A), R^*(p) \tag{14}$$

$$S(p, A) \leftarrow T(p, A), c_S(A), \neg S^-(p), \neg S^+(p, \_) \tag{15}$$

$$S(p, A) \leftarrow S^+(p, A) \tag{16}$$

$$S(p, A) \leftarrow T(p, A), S^*(p) \tag{17}$$

$$T'(p, A) \leftarrow T(p, A), \neg c_R(A), \neg c_S(A), \neg R^*(p), \neg S^*(p) \tag{18}$$

$\gamma_{\mathbf{src}}$:

$$T(p, A) \leftarrow R(p, A) \tag{19}$$

$$T(p, A) \leftarrow S(p, A), \neg R(p, \_) \tag{20}$$

$$T(p, A) \leftarrow T'(p, A) \tag{21}$$

$$R^-(p) \leftarrow S(p, A), \neg R(p, \_), c_R(A) \tag{22}$$

$$R^*(p) \leftarrow R(p, A), \neg c_R(A) \tag{23}$$

$$S^+(p, A) \leftarrow S(p, A), R(p, A'), A \neq A' \tag{24}$$

$$S^-(p) \leftarrow R(p, A), \neg S(p, \_), c_S(A) \tag{25}$$

$$S^*(p) \leftarrow S(p, A), \neg c_S(A), \neg R(p, \_) \tag{26}$$

$$S^*(p) \leftarrow S(p, A), \neg c_S(A), R(p, A) \tag{27}$$

The semantics of all other BiDEL SMOs are defined in a similar way. We do not elaborate on them here.

5. ENSURING BIDIRECTIONALITY

To prove the bidirectional semantics of BiDEL's SMOs we have to consider two cases depending on the SMO's materialization state: data is either (1) materialized on the source side (SMO is virtualized) or (2) materialized on the target side (SMO is materialized). In both cases, data in both schema versions needs to be propagated correctly without any information lost or gained from and to the corresponding data tables. To propagate data at the unmaterialized side, the SMO semantics comes into play.

Let's consider the first case; the data is materialized on the source side. For a correct target-side propagation, the data D_{trg} at the target side has to be mapped by γ_{src} to the data tables and auxiliary tables at the source side (write) and mapped back by γ_{trg} (read) without any loss or gain visible at the target side. This includes that there is no data added to the auxiliary tables on the target site. These conditions have been already defined for symmetric relational lenses [9]—given data at the target side, storing it at the source side, and mapping it back to target should return the identical data but no additional auxiliary data at the target side. For the second case it is vice versa. Formally, an SMO has bidirectional semantics if the following holds:

$$D_{trg} = \gamma_{trg}(\gamma_{src}(D_{trg})) \tag{28}$$

$$D_{src} = \gamma_{src}(\gamma_{trg}(D_{src})) \tag{29}$$

This is shown by simplifying the Datalog rule sets that define the mappings γ_{src} and γ_{trg}. A Datalog rule is a clause of the form $H \leftarrow L_1, \ldots, L_n$ with $n \geq 0$ where H is an atom denoting the rule's head, and $L_1, \ldots, L_n$ are literals, i.e. positive or negative atoms, representing its body. For a given rule r, we use $\mathsf{head}(r)$ to denote its head H and $\mathsf{body}(r)$ to denote its set of body literals $L_1, \ldots, L_n$. If $L \equiv p(v_1, \ldots, v_m)$ with $n \geq 0$ is a literal, we use $\mathsf{pred}(L)$ to refer to the predicate symbol p of L and $\mathsf{vars}(L)$ to denote the set of variables v_i occurring in L. Literals in InVerDa rules do not contain constants. In literals we use capital letters to denote multiple variables v_i. For a set of literals $\mathcal{K}$, $\mathsf{vars}(\mathcal{K})$ denotes $\bigcup_{L \in \mathcal{K}} \mathsf{vars}(L)$. If L is the head of a given rule r, we use $\mathsf{pred}(r)$ to refer to the predicate symbol of L. For a set of rules $\mathcal{R}$, $\mathcal{R}^p$ is defined as $\{r \mid r \in \mathcal{R} \wedge \mathsf{pred}(r) = p\}$. The following lemmas are used for simplifying a given rule set $\mathcal{R}$ into a rule set $\mathcal{R}'$ such that $\mathcal{R}'$ derives the same facts as $\mathcal{R}$.

Lemma 1 (Deduction). *Let $L \equiv p(X, Y)$ be a literal in the body of rule r in the rule set $\mathcal{R}$ with $X = \mathsf{vars}(\mathsf{body}(r) \setminus \{L\}) \cap \mathsf{vars}(L) \neq \emptyset$ and $Y = \mathsf{vars}(L) \setminus X$. For a rule $s \in \mathcal{R}^{\mathsf{pred}(L)}$ let $\mathsf{rn}(s, L)$ be rule s with all variables occurring in the head of s at positions of X variables in L be renamed to match the corresponding X variable and all other variables be renamed to anything not in $\mathsf{vars}(\mathsf{body}(r))$. If L is a positive literal, s can be applied to r to get rule set $r(s)$ of the form $\{\mathsf{head}(r) \leftarrow \mathsf{body}(r) \setminus \{L\} \cup \mathsf{body}(\mathsf{rn}(s, L))\}$. If L is a negative literal, s can be applied to r to get rule set $r(s)$ of the form $\{\mathsf{head}(r) \leftarrow \mathsf{body}(r) \setminus \{L\} \cup \{\neg K\} \mid K \in \mathsf{body}(\mathsf{rn}(s, L))\}$. Accordingly, $\mathcal{R}$ can be simplified by replacing r and all $s \in \mathcal{R}^p$ with all $r(s)$ applications to $\mathcal{R} \setminus (\{r\} \cup \mathcal{R}^p) \cup (\bigcup_{s \in \mathcal{R}^{\mathsf{pred}(L)}} r(s))$.*

Lemma 2 (Empty Predicate). *Let $r \in \mathcal{R}$ be a rule, L be a literal its the body $L \in \mathsf{body}(r)$ and the relation $\mathsf{pred}(L)$ is known to be empty. If L is a positive literal, r can be removed from $\mathcal{R}$. If L is a negative literal, r can be simplified to $\mathsf{head}(r) \leftarrow \mathsf{body}(r) \setminus \{L\}$.*

LEMMA 3 (TAUTOLOGY). *Let $r, s \in \mathcal{R}$ be rules and L and K be literals in the bodies of r and s, respectively, where r and s are identical except for L and K, i.e. $\mathsf{head}(r) = \mathsf{head}(s)$ and $\mathsf{body}(r) \setminus \{L\} = \mathsf{body}(s) \setminus \{K\}$, or can be renamed to be so. If $K \equiv \neg L$, r can be simplified to $\mathsf{head}(r) \leftarrow \mathsf{body}(r) \setminus \{L\}$ and s can be removed from $\mathcal{R}$.*

LEMMA 4 (CONTRADICTION). *Let $r \in \mathcal{R}$ be a rule and L and K be literals in its body $L, K \in \mathsf{body}(r)$. If $K \equiv \neg L$, r can be remove from $\mathcal{R}$.*

Since all tuples have the INVERDA-managed identifier p, the following lemma also applies.

LEMMA 5 (FUNCTIONAL DEPENDENCY). *Let $r \in \mathcal{R}$ be a rule and $q(p, X)$ and $q(p, Y)$ be literals in its body. Since p is a unique identifier, r can be modified to $\mathsf{head}(r) \leftarrow \mathsf{body}(r) \cup \{X = Y\}$.*

Using these lemmas, we can formally show the bidirectionality of all BIDEL SMOs. In this paper, we show bidirectionality for the materialized `PARTITION` SMO in detail. This requires to show that Equation 29 holds. Writing data T_D from the source side to the target-side results in the mapping $\gamma_{trg}(T_D)$. With target-side materialization all source-side auxiliary tables are not needed. Thus, the rule set for $\gamma_{trg}(T_D)$ can be simplified with Lemma 2.

$$R(p, A) \leftarrow T_D(p, A), c_R(A) \tag{30}$$
$$S(p, A) \leftarrow T_D(p, A), c_S(A) \tag{31}$$
$$T'(p, A) \leftarrow T_D(p, A), \neg c_R(A), \neg c_S(A) \tag{32}$$

Reading the source-side data back from R, S, and T' to T adds the rule set γ_{src} (Rule 19–27) to the mapping. Using Lemma 1, the mapping $\gamma_{src}(\gamma_{trg}(T_D))$ simplifies to:

$$T(p, A) \leftarrow T_D(p, A), c_R(A) \tag{33}$$
$$T(p, A) \leftarrow T_D(p, A), c_S(A), \neg T_D(p, A) \tag{34}$$
$$T(p, A) \leftarrow T_D(p, A), c_S(A), \neg c_R(A) \tag{35}$$
$$T(p, A) \leftarrow T_D(p, A), \neg c_S(A), \neg c_R(A) \tag{36}$$
$$R^-(p) \leftarrow T_D(p, A), c_S(A), \neg T_D(p, A), c_R(A) \tag{37}$$
$$R^-(p) \leftarrow T_D(p, A), c_S(A), \neg c_R(A), c_R(A) \tag{38}$$
$$R^*(p) \leftarrow T_D(p, A), c_R(A), \neg c_R(A) \tag{39}$$
$$S^+(p, A) \leftarrow T_D(p, A), c_S(A), T_D(p, A'), c_R(A'), A \neq A' \tag{40}$$
$$S^-(p) \leftarrow T_D(p, A), c_R(A), \neg T_D(p, A), c_S(A) \tag{41}$$
$$S^-(p) \leftarrow T_D(p, A), c_R(A), \neg c_S(A), c_S(A) \tag{42}$$
$$S^*(p) \leftarrow T_D(p, A), c_S(A), \neg c_S(A), \neg T_D(p, \_) \tag{43}$$
$$S^*(p) \leftarrow T_D(p, A), c_S(A), \neg c_S(A), \neg c_R(\_) \tag{44}$$
$$S^*(p) \leftarrow T_D(p, A), c_S(A), \neg c_S(A), T_D(p, A), c_R(A) \tag{45}$$

According to Lemma 4, we can omit Rule 34 as it contains a contradiction. Based on Lemma 3, we omit Rule 36 and simplify Rule 35 to Rule 47 by removing the literal $c_S(A)$. The resulting rules for T

$$T(p, A) \leftarrow T_D(p, A), c_R(A) \tag{46}$$
$$T(p, A) \leftarrow T_D(p, A), \neg c_R(A) \tag{47}$$

can be simplified again with Lemma 3 to

$$T(p, A) \leftarrow T_D(p, A) \quad . \tag{48}$$

For Rule 40, Lemma 5 implies $A = A'$, so this rule can be removed based on Lemma 4. Likewise, the Rules 38–45 have contradicting literals on T_D, c_R, and c_S respectively so that Lemma 4 applies here as well. The resulting rule set clearly shows that the data T_D inserted into D_{src} is mapped by $\gamma_{src}(\gamma_{trg}(D_{src}))$ to the target side and back to D_{src} without any information loss or gain:

$$\gamma_{\mathbf{src}}(\gamma_{\mathbf{trg}}(\mathbf{D_{src}})) : T(p, A) \leftarrow T_D(p, A) \quad \square \tag{49}$$

So, $D_{src} = \gamma_{src}(\gamma_{trg}(D_{src}))$ holds. Remember that the auxiliary tables only exist on the materialized side of the SMO (target in this case). Hence, it is absolutely correct that there are no rules left producing data for the source-side auxiliary.

The same can be done for the other case (Equation 28), where the simplification of $\gamma_{trg}(\gamma_{src}(D_{trg}))$ results in

$$\gamma_{\mathbf{trg}}(\gamma_{\mathbf{src}}(\mathbf{D_{trg}})) : R(p, A) \leftarrow R_D(p, A) \tag{50}$$
$$S(p, A) \leftarrow S_D(p, A) \tag{51}$$

The shown formal validation can be applied to the remaining BIDEL SMOs, as well. The BIDEL SMO definitions ensure that given data at any schema version V_n that is propagated and stored at a direct predecessor V_{n-1} or direct successor schema version V_{n+1} can always be read completely and correctly in V_n. To the best of our knowledge, we are the first to design a set of powerful SMOs and validate their bidirectionality according to the criteria of symmetric relational lenses.

Write operations. This bidirectionality also holds after write operations: When updating a not-materialized schema version, this update is propagated to the materialized schema in a way that it is correctly reflected when reading the updated data again. Given a materialized SMO, we apply a write operation $\Delta_{src}(D_{src})$ to given data on the source side. Initially, we store D_{src} at the target side using $D_{trg} = \gamma_{trg}(D_{src})$. To write at the source side, we have to temporarily map the data D_{trg} back to the source, apply the write Δ_{src}, and map the updated data back to target with $D'_{trg} = \gamma_{trg}(\Delta_{src}(\gamma_{src}(D_{trg})))$. Reading the updated target $\gamma_{src}(D'_{trg})$ has to be equal to applying the write operation $\Delta_{src}(D_{src})$ directly on the source side.

$$\Delta_{src}(D_{src}) = \gamma_{src}(\gamma_{trg}(\Delta_{src}(\gamma_{src}(\gamma_{trg}(D_{src})))))) \tag{52}$$

We have already shown that $D = \gamma_{src}(\gamma_{trg}(D))$ holds for any data D at the target side, so that Equation 52 reduces to $\Delta_{src}(D_{src}) = \Delta_{src}(D_{src})$. Hence writes are correctly propagated through the SMOs. The same holds vice versa for writing at the target-side of virtualized SMOs:

$$\Delta_{trg}(D_{trg}) = \gamma_{trg}(\gamma_{src}(\Delta_{trg}(\gamma_{trg}(\gamma_{src}(D_{trg})))))) \tag{53}$$

Chains of SMOs. Further, the bidirectionality of BIDEL SMOs also holds for chains of SMOs: $smo_1, \ldots smo_n$, where $\gamma^i_{src/trg}$ is the respective mapping of smo_i. Analogously to symmetric relational lenses [9], there are no side-effects between multiple BIDEL SMOs. So, BIDEL's bidirectionality is also guaranteed along chains of SMOs:

$$D_{trg} = \gamma^n_{trg}(\ldots \gamma^1_{trg}(\gamma^1_{src}(\ldots \gamma^n_{src}(D_{trg})))) \tag{54}$$
$$D_{src} = \gamma^1_{src}(\ldots \gamma^n_{src}(\gamma^n_{trg}(\ldots \gamma^1_{trg}(D_{src})))) \tag{55}$$

With the formal guarantee of bidirectionality—also along chains of SMOs and for write operations—we have laid a solid formal foundation for INVERDA's delta code generation.

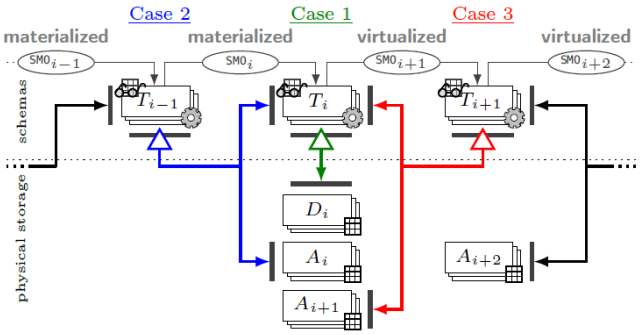

Figure 6: Three different cases in delta code generation.

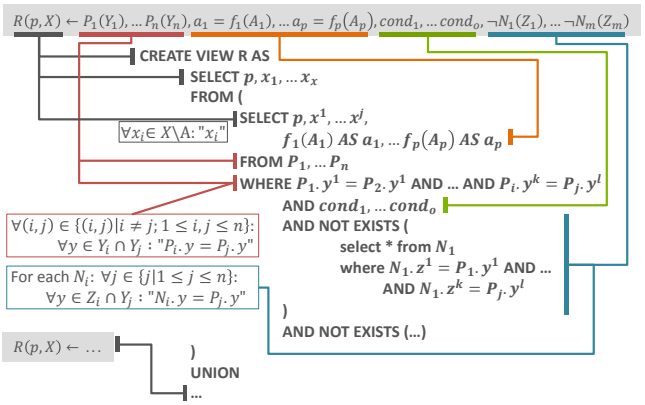
$R(p,X) \leftarrow P_1(Y_1), \ldots P_n(Y_n), a_1 = f_1(A_1), \ldots a_p = f_p(A_p), cond_1, \ldots cond_o, \neg N_1(Z_1), \ldots \neg N_m(Z_m)$
CREATE VIEW R AS
SELECT $p, x_1, \ldots x_x$
FROM (
SELECT $p, x^1, \ldots x^j,$
$f_1(A_1)$ AS $a_1, \ldots f_p(A_p)$ AS a_p
$\forall x_i \in X \backslash A$: "$x_i$"
FROM $P_1, \ldots P_n$
WHERE $P_1.y^1 = P_2.y^1$ AND ... AND $P_i.y^k = P_j.y^l$
AND $cond_1, \ldots cond_o$
$\forall (i,j) \in \{(i,j) | i \neq j; 1 \leq i,j \leq n\}$:
$\forall y \in Y_i \cap Y_j$: "$P_i.y = P_j.y$"
AND NOT EXISTS (
select * from N_1
where $N_1.z^1 = P_1.y^1$ AND ...
AND $N_1.z^k = P_j.y^l$
For each N_i: $\forall j \in \{j | 1 \leq j \leq n\}$:
$\forall y \in Z_i \cap Y_j$: "$N_i.y = P_j.y$"
)
AND NOT EXISTS (...)
$R(p,X) \leftarrow \ldots$
)
UNION
...

Figure 7: SQL generation from Datalog rules.

6. DELTA CODE GENERATION

To make a schema version available, InVerDa translates the γ_{src} and γ_{trg} mapping functions of SMOs into delta code—specifically views and triggers. Views implement delta code for read operations; triggers implement delta code for write operations. In a genealogy of schema versions, a single table's version can be involved in SMO instances as target, source, or both. The delta code for a specific table version depends on the materialization state of the table's adjacent SMOs.

Consider the schema genealogy shown in Figure 6. The two subsequent SMO instances, $i-1$ and i are set to store their data at the target side (materialized), while the two subsequent SMO instances, $i+1$ and $i+2$ are set to source-side materialization (virtualized). Without loss of generality, three cases for delta code generation can be distinguished. *Case 1:* The incoming SMO is materialized and all outgoing SMOs are virtualized. The data of a table version T_i is physically stored in the same version. The view and the trigger that implement T_i directly propagate reads and writes to data table D_i without any schema translation necessary. *Case 2:* The incoming SMO is materialized and one outgoing SMO is also materialized. The data of a table version T_{i-1} is physically stored upstream the schema genealogy, i.e. in newer versions. The view and the trigger that implement T_{i-1} propagate along the materialized outgoing `SMO`$_i$ to views and auxiliary tables of schema version i. Schema translation happens according to γ_{src} (read data) and γ_{trg} (written data) of `SMO`$_i$. *Case 3:* The incoming SMO and the outgoing SMOs are virtualized. The data of a table version T_{i+1} is physically stored downstream the schema genealogy. The view and the trigger that implement T_{i+1} propagate back along the virtualized incoming `SMO`$_{i+1}$ to views and auxiliary tables of version i. Schema translation happens according to γ_{trg} (read data) and γ_{src} (written data) of `SMO`$_{i+1}$. In Case 1 delta code generation is trivial. In the other two cases, InVerDa essentially translates the Datalog rules defining the relevant mapping functions into view and trigger definitions.

For reads, Figure 7 illustrates the general pattern of the translation of Datalog rules to a view definition. As a single table can be derived by multiple rules, e.g. Rule 19–21, a view is a union of subqueries each representing one of the rules. For each subquery InVerDa lists all attributes of the rule head in the select clause. Within a nested subselect these attributes are either projected from the respective table version or derived by a function (e.g. values for an added column). All positive literals referring to other table versions or auxiliary tables are listed in the from clause. Further, InVerDa adds for all attributes occurring in multiple positive literals respective join conditions to the where clause. Finally, conditions, such as $c_S(X)$, and negative literals, which InVerDa adds as a `NOT EXISTS(<subselect for the literal>)` condition, complete the where-clause.

For writes, InVerDa creates three triggers on each table version—for inserts, deletes, and updates respectively. To not recompute the whole data of the materialized side whenever a write operation is executed at the not-materialized side of an SMO, InVerDa adopts an update propagation technique for Datalog rules [2] that results in minimal write operations. For instance an insert operation $\Delta_T^+(p,A)$ on the table version T propagated back to the source side of a materialized `PARTITION` SMO results in the following update rules:

$$\Delta_R^+(p,A) \leftarrow \Delta_T^+(p,A), \texttt{new}\ c_R(A), \texttt{old}\ \neg R(p,A) \quad (56)$$

$$\Delta_S^+(p,A) \leftarrow \Delta_T^+(p,A), \texttt{new}\ c_S(A), \texttt{old}\ \neg S(p,A) \quad (57)$$

$$\Delta_{T'}^+(p,A) \leftarrow \Delta_T^+(p,A), \texttt{new}\ \neg c_R(A), \neg c_S(A), \texttt{old}\ \neg T'(p,A) \quad (58)$$

The propagation can handle multiple write operations at the same time and distinguishes between `old` and `new` data, which represents the state before and after applying other write operations. The derived update rules match the intuitive expectations: The inserted tuple is propagated to R or to S or to T' given it satisfies either c_R or c_S or none of them. The additional conditions on the `old` literals ensures minimality by checking whether the tuple already exists in the respective table. To generate the trigger code from the update rule resulting from the propagation, InVerDa applies essentially the same algorithm as for view generation.

Writes performed by a trigger on another table version trigger the propagation further along the schema genealogy the other table versions as long as the respective update rules deduce write operations, i.e. as long as some data is physically stored with a tables version either in the data table or in auxiliary tables. With the generated delta code, InVerDa propagates writes on any schema version to every other co-existing schema version in a schema genealogy.

M	P
$\emptyset$	{Task-0}
{PARTITION}	{Task-0}
{PARTITION, DROP COLUMN}	{Todo-1}
{DECOMPOSE}	{Task-1, Author-0}
{DECOMPOSE, RENAME}	{Task-1, Author-1}

Table 1: Possible materialization schemas and the corresponding physical table schema in the TasKy example.

InVerDa	Initially	Evolution	Migration
Lines of Code	10	4	1
Statements	10	4	1
Characters	1230	229	31
SQL (Ratio)	**Initially**	**Evolution**	**Migration**
Lines of Code	10 (×1.00)	87 (×21.75)	264 (×264.00)
Statements	10 (×1.00)	55 (×13.75)	140 (×140.00)
Characters	1275 (×1.04)	3636 (×15.88)	9911 (×319.71)

Table 2: Ratio between SQL and InVerDa delta code.

7. MIGRATION PROCEDURE

The materialization states of all SMO instances in a schema genealogy form the materialization schema. The materialization schema determines the physical table schema, i.e. which table versions are directly stored in physical storage. For the TasKy example—Figure 1—this entails five different possible materialization schemas M, each implying a different physical table schema P as shown in Table 1.

The materialization schema has huge impact on the performance of a given workload. Workload changes such as increased usage of newer schema versions demand adaptations of the materialization schema. InVerDa facilitates such an adaptation with a foolproof migration command. The migration command allows moving data non-redundantly along the schema genealogy to those table versions where the given workload causes the least overhead for propagating reads and write. Initially, all SMOs except of the create table SMOs are virtualized, i.e. only initially created table versions are in the physical table schema. A new materialization schema is derived from a given valid materialization schema by changing the materialization state of selected SMO instances. Formally, two conditions must hold for a materialization schema to be valid. For an SMO s, we denote the source table versions as $\texttt{src}(s)$. For each table version t we denote the incoming SMO with $\texttt{in}(t)$ and the set of outgoing SMOs with $\texttt{out}(t)$. A materialization is valid iff:

$$\forall s \in M\ \forall t \in \texttt{src}(s)\ \ (\texttt{in}(t) \in M) \tag{59}$$

$$\forall s \in M\ \forall t \in \texttt{src}(s)\ \nexists o \in (\texttt{out}(t) \setminus \{s\})\ \ (o \in M) \tag{60}$$

The first condition ensures that all source table versions are in the materialization schema. The second condition ensures that no source table version is already taken by another materialized SMO.

In the migration command, the DBA lists the table versions that should be materialized. For instance:

```
MATERIALIZE 'TasKy2.task', 'TasKy2.author';
```

InVerDa determines with the schema version catalog the corresponding materialization schema and checks, whether it is valid according to the conditions above. If valid, InVerDa automatically creates the new physical tables including auxiliary tables in the physical storage, migrates the data, regenerates all necessary delta code, and deletes all old physical tables. The actual data migration relies on the same SQL generation routines as used for view generation. From a user perspective all schema versions still behave the same after a migration. However, any data access is now propagated to the new physical table schema resulting in a better performance for schema version evolution-wise close to this new physical schema.

8. EVALUATION

InVerDa brings great advantages for software systems by decoupling the different goals of different stakeholders. Users can continuously access all schema versions, while developers can focus on the actual continuous implementation of the software without caring about former versions. Above all, the DBA can change the physical table schema of the database to optimize the overall performance without restricting the availability of the co-existing schema versions or invalidating the developers' code. In Section 8.1, we show how InVerDa reduces the length and complexity of the code to be written by the developer and thereby yields more robust and maintainable solutions. InVerDa automatically generates the delta code based on the discussed Datalog rules. We measure the overhead of accessing data through InVerDa's delta code compared to a handwritten SQL implementation of co-existing schema versions in Section 8.2 and show that it is reasonable. In Section 8.3, we show that the possibility to easily adapt the physical table schema to a changed workload outweighs the small overhead of InVerDa's automatically generated delta code. Materializing the data according to the most accessed version speeds up the data access significantly.

For the measurements, we use both our TasKy example and an evolved TPC-H example. From the original TPC-H [16] schema (denoted $TPCH_0$), we derive two new versions. For schema version $TPCH_1$, we denormalize the schema by joining the Nation with Region. We also drop the regions' comment since nation comes with its own comment column:

```
CREATE SCHEMA VERSION TPCH1 FROM TPCH0 WITH
DROP COLUMN r_comment FROM Region
     DEFAULT 'comment_'+rand();
OUTER JOIN TABLE Nation, Region INTO Country
     ON FOREIGN KEY n_regionkey;
```

For schema version $TPCH_2$, we rename the Part.p_mfgr into p_manufacturer and partition the Part into normal and oversize parts. Additionally we add the column p_surcharge to the oversize parts:

```
CREATE SCHEMA VERSION TPCH2 FROM TPCH0 WITH
RENAME COLUMN p_mfgr IN Part TO p_manufacturer;
PARTITION TABLE Part INTO
               OversizePart WITH p_size>=40,
               NormalsizePart WITH p_size<40;
ADD COLUMN p_surcharge INT AS 20 TO OversizePart;
```

InVerDa's current implementation uses p as a system managed primary key for each table. We add indices on those columns that are primary keys in TPC-H to obtain a comparable performance. We measure single thread performance of a PostgreSQL 9.4 database with co-existing schema versions on a Core i7 machine with 2,4GHz and 8GB memory.

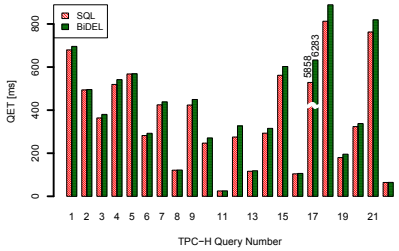

Figure 8: Benchmark $TPCH_0$ with materialized $TPCH_1$.

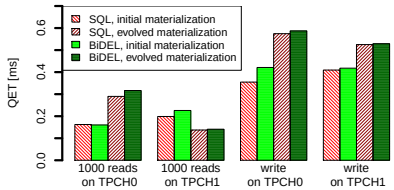

Figure 9: QET depending on the materialization.

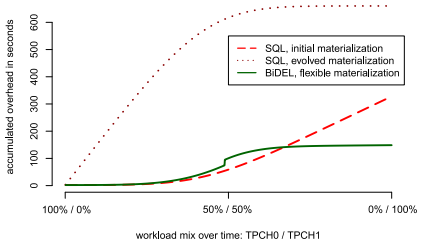

Figure 10: Performance gain with flexible materialization.

8.1 Simplicity and Robustness

Most importantly, we show that InVerDa unburdens developers by rendering the expensive and error-prone task of manually writing delta code unnecessary. To show this, we implement the first evolution to $TPCH_1$ with handwritten and hand-optimized SQL and compare this code to the equivalent BiDEL statements. We did this for (1) creating the initial $TPCH_0$ schema, (2) creating the additional schema version $TPCH_1$ by creating the respective views and triggers, and (3) migrating the physical table schema to $TPCH_1$ and adapting all existing delta code.

The handwritten delta code is much longer and much more complex than achieving the same goal with BiDEL. Table 2 shows the lines of code (LOC) required with SQL and BiDEL, respectively, as well as the ratio between these values. As there is no general coding style for SQL, LOC is a rather vague measure. We also include the objective number of statements and number of characters (consecutive white-space characters counted as one) to get a clear picture and use the number of characters per statement as indicator for the code's complexity. Obviously, creating the initial schema is equally complex for both approaches. However, evolving to the new schema version $TPCH_1$ and migrating the data accordingly requires 87 and 264 LOC respectively, while we can express the same with 4 and 1 lines with BiDEL. Moreover, the SQL code is also more complex, as can be seen by the average number of character per statement. While BiDEL is working exclusively on the external schema versions, with handwritten SQL developers also have to manage auxiliary tables to prevent information loss, propagated write operations to all versions correctly, etc.

The automated delta code generation does not only eliminate the error-prone and expensive manual implementation, but it is also reasonably fast. Creating the initial $TPCH_0$ took 145 ms on our test system. The evolution to $TPCH_1$, which includes two SMOs, requires 184 ms for both the generation and execution of the evolution script. The same took 213 ms for $TPCH_2$ which contains already three SMOs. Please note that the complexity of generating and executing evolution scripts depends linearly on the number of SMOs within an evolution N and the number of untouched table versions M. The complexity is $O(N+M)$, since we generate the delta code for each SMO locally and exclusively work on the neighboring table versions. This principle protects against additional complexity in longer chains of SMOs. The same holds for the complexity of executing migration scripts. It is $O(N)$ since InVerDa merely moves the data and updates the delta code for the materialized SMOs stepwise.

8.2 Overhead of Generated Delta Code

InVerDa's generated delta code is derived from the presented Datalog rules and aims at a general and solid solution. So far, our focus is on the correct propagation of data access on multiple co-existing schema versions. We assume the database's optimizer to find a fast execution plan, however, there will be an overhead of InVerDa compared to hand-optimized SQL.

The essential aspect of our scenario are co-existing schema versions, hence our handwritten baseline with SQL also provides the co-existing schema versions. The hand-optimized delta code causes slightly less overhead than the generated one but requires significantly more LOC and makes adapting the physical table schema to the current workload a tough job. To show this, we execute the unchanged TPC-H benchmark on $TPCH_0$ with $TPCH_1$ being materialized. To understand the performance overhead in detail, we conduct a micro benchmark on the affected tables.

TPC-H benchmark. Figure 8 shows the measured read performance for the TPC-H (SF=0.1) benchmark on the original schema version $TPCH_0$, with $TPCH_1$ being materialized. Even these rather complex queries do only lead to a small overhead of InVerDa's generated delta code compared to handwritten SQL. Keeping in mind the difference in length and complexity of the code (264 x LOC for the migration), a performance overhead of 5 % in average is more than reasonable for most users.

Micro benchmark. Figure 9 shows the overhead when accessing only the affected Nation table in $TPCH_0$ and the Country table in $TPCH_1$ with 10^7 tuples. The read workload queries the respective table without any projection or selection. There are two major insights. First, the overhead of the InVerDa generated delta code is negligibly small compared to the handwritten SQL. Second, the materialization significantly influences the actual performance. Reading the data in the materialized version is up to twice as fast as accessing it from the respective other version in this scenario. The write workload inserts new tuples to Nation and Country respectively. Again, the overhead of InVerDa compared to SQL is reasonable small. However, the initial materialization is always faster because the evolved materialization requires to manage an additional auxiliary table for the foreign key relationship. A DBA can optimize the overall performance for a given workload by adapting the materialization, which is itself a very simple task with InVerDa. An advisor tool supporting the optimization task is very well imaginable, but out of scope for this paper.

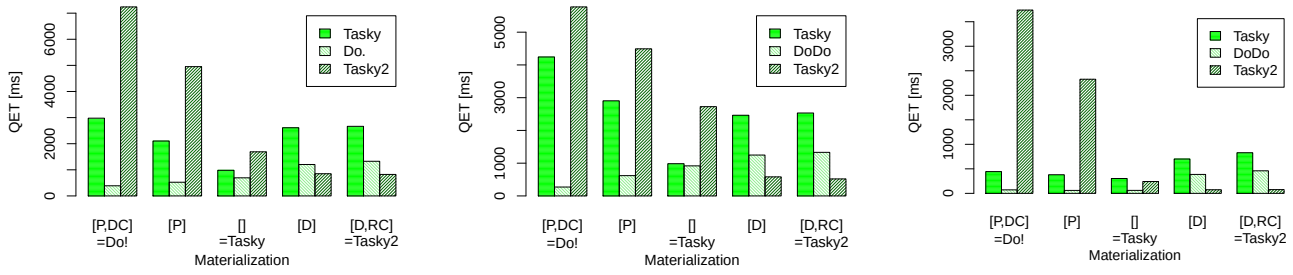

(a) TasKy. (b) TasKy Read. (c) TasKy Write.

Figure 11: Materialization of TasKy versions.

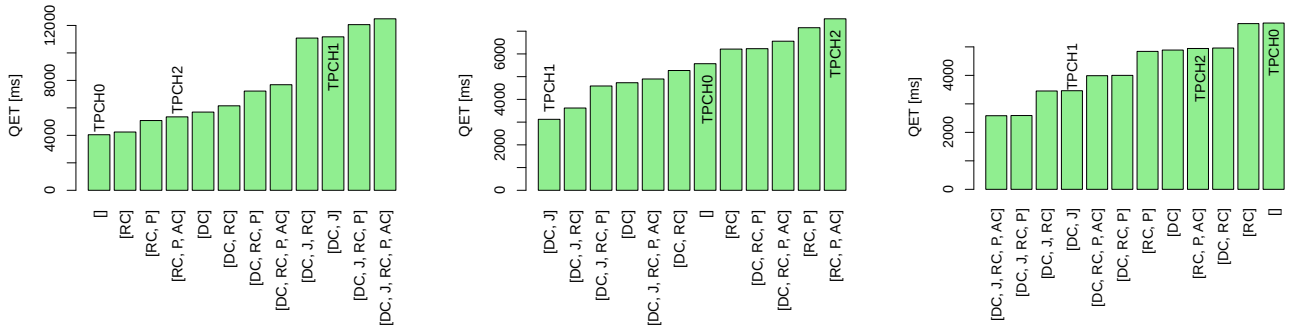

(d) Access $TPCH_0$. (e) Access $TPCH_1$. (f) Access $TPCH_2$.

Figure 12: Materialization of TPC-H versions.

8.3 Benefit of Flexible Materialization

Adapting the physical table schema to the current workload is hard with handwritten SQL since all delta has to be rewritten, but almost for free with InVerDa (1 LOC instead of 264 in our example). Let's assume a development team spares the effort for rewriting delta code and works with a fixed materialization. Figure 10 shows the accumulated propagation overhead for handwritten SQL with the two fixed materialization and for InVerDa with an adapted materialization. Assume, over time the workload changes from 0 % access to $TPCH_1$ and 100 % to $TPCH_0$ to the opposite 100 % and 0 % according to the Technology Adoption Life Cycle. The adoption is divided into 1000 time slices where 1000 queries are executed respectively. The workload mixes 50 % reads, 20 % inserts, 20 % updates, and 10 % deletes. As soon as the evolved materialization provides the higher performance for the current workload mix, we instruct InVerDa to change the materialization. As can be seen, InVerDa allows significantly better performance in total—including migration cost—than one of two fixed materialization. So, InVerDa allows to freely choose the materialization to match the current workload without additional development efforts.

The DBA can choose between multiple materialization schemas. The number of valid materialization schemas greatly depends on the actual structure of the evolution. The lower bound is a linear sequence of depending SMOs, e.g. one table with N `ADD COLUMN` SMOs has N valid materializations. The upper bound are N independent SMOs, each evolving another table, with 2^N valid materializations. Specifically, the TPC-H example has twelve valid materializations while TasKy example has five.

Let's first look at TasKy example. Figure 11 shows the data access performance on the three schema versions given a particular materialization schema. The materialization schemas are represented as the lists of SMOs that are materialized. We use abbreviations of the SMOs: e.g. $[D, RC]$ on the very right corresponds to schema version TasKy2 since both the decompose (D) and the rename column (RC) SMO are materialized. The initial materialization is in the middle, while the materialization according to Do! is on the very left. The workload mixes 50 % reads, 20 % inserts, 20 % updates, 10 % deletes in Figure 11a, 100 % reads in Figure 11b, and 100 % inserts in Figure 11c on the depicted schema versions. Again, the measurements show that accesses to each schema version are fastest when exactly the table version of that schema are materialized, i.e when the materialization schema fits the schema version. However, there are differences in the actual overhead, so the globally optimal materialization depends on the workload distribution among the schema version. E.g. writing to TasKy2 is 49 times faster when the physical table schema matches TasKy2. This gain increases with every SMO, so for longer evolutions with more SMOs, the benefit is even higher in the order of magnitudes.

For the TPC-H example results are shown Figure 12. The subfigures depict the performance for accessing each schema version with the same workload mix (50 % read, 20 % insert, 20 % update, 10 % delete). Each bar represent another materialization schema. Again, each schema version performs best with the corresponding materialization schema with one interesting exception. The evolutions from $TPCH_0$ to $TPCH_1$ and from $TPCH_0$ to $TPCH_2$ are independent as they work on different tables. For accessing $TPCH_2$ for instance it

is beneficial to also materialize the SMOs for $TPCH_1$, as write operations are propagated to all table versions and the join SMO introduces a lower overhead when materialized, as already see in Figure 9. In sum, the materialization greatly determines the access performance. InVerDa enables the DBA to easily adapt the materialization schema to a changing workload to significantly speed up query processing without hitting other stakeholder's interests.

9. RELATED WORK

Both the evolution of databases [12] and co-existing schema versions [14] are well recognized in the database research community. For database evolution, there are approaches to increase comfort and efficiency, for instance by defining a schema evolution aware query language [13] or by providing a general framework to describe database evolution in the context of evolving applications [7]. Currently, PRISM [4] appears to provide the most advanced database evolution tool including an SMO-based DEL. PRISM was first introduced in 2008 and focused on the plain database evolution [5]. Later, the authors extended it to PRISM++, which includes the modification of constraints and update rewriting [6].

InVerDa extends database evolution to support co-existing schema versions, which is located in the research field of bidirectional transformations [15]. Particularly, symmetric relational lenses lay a foundation to describe the mapping of read and write accesses along a bidirectional mapping [9]. For InVerDa, we adapt this idea to bidirectional SMOs. Another extension of PRISM++ takes a first step towards co-existing schema versions by answering queries on former schema versions according to the current data [11], however, InVerDa also covers write operations on those former schema versions. There are multiple systems also taking this step, however, the DELs are usually rather limited or work on different meta models like data warehouses [1]. The ScaDaVer system [17] allows additive and subtractive SMOs on the relational model, which simplifies bidirectionality and hence is a great starting point towards more powerful DELs. BiDEL also covers restructuring SMOs and is based on established DELs [4, 8]. To the best of our knowledge, we are the first to realize end-to-end support for co-existing schema versions based on such powerful DELs.

10. CONCLUSIONS

Current DBMSs do not support co-existing schema versions properly, forcing developers to manually write complex and error-prone delta code, which propagates read/write accesses between schema versions. Moreover, this delta code needs to be adapted manually whenever the DBA decides to change the physical table schema. InVerDa greatly simplifies creating and maintaining co-existing schema versions for developers, while the DBA can freely change the physical table schema. For this sake, we have introduced BiDEL, a bidirectional database evolution language that is simple and intuitive, but carries enough information to generate all the delta code automatically. We have formally validated BiDEL's bidirectionality making it a sound and robust basis for InVerDa. In the evaluation, we have shown that BiDEL scripts are significantly shorter than handwritten SQL scripts (264 x). The performance overhead caused by the automatically generated delta code is very low but the freedom to easily change the physical table schema is highly valuable: we can greatly speed up query processing when the physical table schema matches the most accessed version. In sum, InVerDa finally enables agile—but also robust and maintainable—software development for information systems. Future research will aim at (1) even more efficient physical table schemas e.g. by allowing (partial) redundancy, (2) a self-managed DBMS continuously adapting the physical table schema to the current workload, and (3) optimized delta code within a database system instead of triggers.

11. ACKNOWLEDGMENTS

This work is funded by the German Research Foundation (DFG) within the RoSI Research Training Group (1907).

12. REFERENCES

[1] M. Arora and A. Gosain. Schema Evolution for Data Warehouse: A Survey. *IJCA*, 22(5):6–14, 2011.

[2] A. Behrend, U. Griefahn, H. Voigt, and P. Schmiegelt. Optimizing continuous queries using update propagation with varying granularities. In *SSDBM'15*, pages 1–12, 2015.

[3] M. L. Brodie and J. T. Liu. Keynote: The Power and Limits of Relational Technology In the Age of Information Ecosystems. In *OTM'10*, pages 2–3, 2010.

[4] C. Curino, H. J. Moon, A. Deutsch, and C. Zaniolo. Automating the database schema evolution process. *VLDB J*, 22(1):73–98, 2013.

[5] C. Curino, H. J. Moon, and C. Zaniolo. Graceful database schema evolution: the PRISM workbench. *PVLDB*, 1(1):761–772, 2008.

[6] C. Curino and C. Zaniolo. Update rewriting and integrity constraint maintenance in a schema evolution support system: PRISM. *PVLDB*, 4(2):117–128, 2010.

[7] E. Domínguez, J. Lloret, Á. L. Rubio, and M. a. Zapata. MeDEA: A database evolution architecture with traceability. *DKE*, 65(3):419–441, 2008.

[8] K. Herrmann, H. Voigt, A. Behrend, and W. Lehner. CoDEL - A Relationally Complete Language for Database Evolution. In *ADBIS'15*, pages 63–76, 2015.

[9] M. Hofmann, B. Pierce, and D. Wagner. Symmetric lenses. *POPL'11*, 46(1), 2011.

[10] P. Howard. Data migration report, 2011.

[11] H. J. Moon, C. Curino, M. Ham, and C. Zaniolo. PRIMA - archiving and querying historical data with evolving schemas. In *SIGMOD'09*, pages 1019–1022, 2009.

[12] E. Rahm and P. A. Bernstein. An online bibliography on schema evolution. *SIGMOD Rec*, 35(4):30–31, 2006.

[13] J. F. Roddick. SQL/SE - A Query Language Extension for Databases Supporting Schema Evolution. *SIGMOD Rec*, 21(2):10–16, 1992.

[14] J. F. Roddick. A survey of schema versioning issues for database systems. *IST*, 37(7):383–393, 1995.

[15] J. F. Terwilliger, A. Cleve, and C. Curino. How Clean Is Your Sandbox? *LNCS*, 7307:1–23, 2012.

[16] Transaction Processing Performance Council (TPC). TPC-H Benchmark. *http://www.tpc.org/tpcc/*.

[17] B. Wall and R. Angryk. Minimal data sets vs. synchronized data copies in a schema and data versioning system. In *PIKM'11*, page 67, 2011.